DONGKEUN KIM

1418 A.V.Williams Bldg.
ECE Department, University of Maryland
College Park, MD 20742 USA

Phone: (301) 405-2068
E-mail: dongkeun@eng.umd.edu
URL: http://www.ece.umd.edu/~dongkeun

Objective

A research and development position to design high performance computer systems and develop optimizing compilers for multi-threaded processors or embedded systems

Education

University of Maryland, College Park, MD
Ph.D., Electrical and Computer Engineering, Fall 2004 (Expected)
M.S., Electrical and Computer Engineering, December 2003

Seoul National University, Seoul, Korea
B.S., Electrical Engineering, February 1999

Working Experience

Microarchitecture Research Lab, Intel Corporation, Santa Clara, CA
Intern (Supervisor: Dr. John Shen), March 2003 ~ May 2004

Department of ECE, University of Maryland, College Park, MD
Research Assistant (Thesis Advisor: Dr. Donald Yeung), June 2000 ~ February 2003

Department of ECE, University of Maryland, College Park, MD
Teaching Assistant, ENEE446: Digital Computer Design, January 2000 ~ May 2000

Publications

Articles in Refereed Journals

1 Dongkeun Kim and Donald Yeung. A Study of Source-Level Compiler Algorithms for Automatic Construction of Pre-Execution Code. *ACM Transactions on Computer Systems (TOCS)*. Vol. 22, No. 3, August 2004

2 Seungryul Choi, Nicholas Kohout, Sumit Pamnani, Dongkeun Kim, and Donald Yeung. A General Framework for Prefetch Scheduling in Linked Data Structures and its Application to Multi-Chain Prefetching. *ACM Transactions on Computer Systems (TOCS)*. Vol. 22, No. 2. May 2004

Articles in Refereed Conferences

1 Nicholas Kohout, Seungryul Choi, Dongkeun Kim, and Donald Yeung. Multi-Chain Prefetching: Effective Exploitation of Inter-Chain Memory Parallelism for Pointer-Chasing Codes. *Proceedings of the 10th Annual International Conference on Parallel Architectures and Compilation Techniques (PACT'01)*. Barcelona, Spain, September 2001

2 Dongkeun Kim and Donald Yeung. Design and Evaluation of Compiler Algorithms for Pre-Execution. *Proceedings of the 10th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-X)*. San Jose, CA, October 2002

3 Steve Shih-wei Liao, Xinmin Tian, Perry Wang, Dongkeun Kim, Juan del Cuvillo, Hong Wang, and John Shen. AutoHelper: Profile-Guided Generation of Helper Threads. *The First Intel Programming Technology Conference (IPTC'03) on Dynamic Compilation and Profiling Guided Optimizations*. Hillsboro, OR, November 2003

4 Dongkeun Kim, Juan del Cuvillo, Steve Shih-wei Liao, Perry Wang, Xinmin Tian, Hong Wang, and John Shen. EmonLite: User-Level Library Routines for Dynamic Performance Monitoring with Low Profiling Overhead. *The First Intel Programming Technology Conference (IPTC'03) on Dynamic Compilation and Profile-guided Optimizations*. Hillsboro, OR, November 2003

5 Dongkeun Kim, Shih-wei Liao, Perry Wang, Juan del Cuvillo, Xinmin Tian, Xiang Zou, Hong Wang, Donald Yeung, Milind Girkar, and John Shen. Physical Experimentation with Prefetching Helper Threads on Intel's Hyper-Threaded Processors. *Proceedings of the Second Annual IEEE/ACM International Symposium on Code Generation and Optimization (CGO2004) with special emphasis on feedback-directed and runtime optimization*. Palo Alto, CA, March 2004 – Best paper award nominee

6 Perry H. Wang, Jamison D. Collins, Hong Wang, Dongkeun Kim, Bill Greene, Kai-Ming Chan, Aamir B. Yunus, Terry Sych, and John P. Shen. Helper Threads via Virtual Multi-Threading on An Experimental Itanium 2 Machine. *Proceedings of the 11th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-XI).* Boston, MA, October 2004

Technical Reports

1 Dongkeun Kim and Donald Yeung. Using Program Slicing to Drive Pre-Execution on Simultaneous Multithreading Processors. University of Maryland Institute for Advanced Computer Studies Technical Report, UMIACS-TR-2001-49, June 2001

Awards

1 Division Recognition Award, Microprocessor Research Labs, Intel Corporation, May 2003

2 MRL Award of Excellence, "TSLRP Demo Team", Intel Corporation, May 2003

Talks

1 Design and Evaluation of Compiler Algorithms for Pre-Execution. Graduate Student Seminar, ECE Department, University of Maryland, College Park, MD, October 2002

2 Design and Evaluation of Compiler Algorithms for Pre-Execution. ASPLOS-X, San Jose, CA, October 2002

3 EmonLite: User-Level Library Routines for Dynamic Performance Monitoring with Low Profiling Overhead. IPTC'03, Intel Corporation, Hillsboro, OR, November 2003

4 Physical Experimentation with Prefetching Helper Threads on Intel's Hyper-Threaded Processors. Microprocessor Technology Lab., Intel Corporation, Santa Clara, CA, March 2004

5 Physical Experimentation with Prefetching Helper Threads on Intel's Hyper-Threaded Processors. CGO2004, Palo Alto, CA, March 2004

6 Physical Experimentation with Prefetching Helper Threads on Intel's Hyper-Threaded Processors & Dynamic Helper Threading. Microarchitecture Research Lab Face-to-Face Event, Intel Corporation, Austin, TX, May 2004

7 Compiler-Based Pre-Execution: Design and Evaluation. System LSI Division, Samsung Electronics Co. Ltd., Yong-In, Korea, July 2004

Patents

1 Method and Apparatus of Compiler-Creating Helper Threads on Multithreading Machines, Liao et. al., Intel Corporation, filed for US patent in Q3 2003, pending

2 Method and Apparatus for Static and Dynamic Thread Management for Processors Supporting Virtual Multithreading, Hoflehner et. al., Intel Corporation, filed for US patent in Q3 2003, pending

3 A Compiler Transformation for Reducing Memory Latency through Helper Threads on Simultaneous Multithreading Architectures, Tian et. al., Intel Coporation, filed for US patent in Q3 2003, pending

Professional Service

Referee for International Conference on Parallel Processing (ICPP-04), Montreal, Canada, August 2004

Coursework

Parallel Processing Computer Architectures, Performance Evaluation of Computer Systems
Design of Distributed Computer Systems, Computer-Aided Design of Digital Systems
Advanced Topics in Microarchitecture, Digital Computer Design, Operating System
Electrical Network Theory, Integrated and Microwave Electronics, Applied Stochastic Processes

Research Description

Due to speed gap between processor and memory system, the memory latency becomes one of the major performance bottlenecks and, moreover, it keeps increasing exponentially. To tolerate the ever-increasing memory latency in high performance computer systems, there have been plethora of proposals from academia and industry. Especially, my research focus is on *Data Prefetching* techniques on multi-threaded processors such as SMT (Simultaneous Multi-Treading) or CMP (Chip Multi-Processor). Currently, I am working on *helper threading* idea, in which a number of prefetch threads run in front of the main computation and trigger cache misses earlier, thereby relieving the main thread from memory stalls. I built a source-level compiler system to construct

effective helper threads. The future works include dynamic control of helper threads via run-time performance monitoring and investigation of the effectiveness of helper threading in the context of multi-programmed workloads.

Skills

Language: C, C++, Assembly (x86, IA-64, MIPS), VHDL, Perl, BASIC

OS: Sun Solaris, Microsoft Windows, Microsoft DOS

Tool: SimpleScalar Tool Set, Unravel, Eel, Shade, Microsoft Office, Sun StarOffice, GNU tools

Compiler: Stanford University Intermediate Format (SUIF) framework, Microsoft Visual Studio Intel Research Compiler Infrastructure for IA-32 and IA-64, GCC (GNU C Compiler)

Citizenship Republic of Korea

References Available upon request